Visual Learning for Real-World Interaction

Trevor Darrell
University of California, Berkeley, USA

Abstract

Contemporary vision research focuses on recognition challenges and data derived from media available on the Internet. While the large-scale datasets this has enabled have led to great progress, methods tuned to these challenges can surprisingly underperform on many real-world problems, and miss opportunities afforded by situated sensing strategies. In this talk I'll present recent results which leverage environment and domain constraints for large-scale recognition: I'll review new methods for domain adaptation, schemes for learning for complex fine-grained categories from limited training data using pose-normalized descriptors, and techniques for performing time-optimal "anytime recognition" with a given computational budget.

Keywords: Domain adaptation, Fine-grained recognition, and Budgeted recognition.